

What is claimed is:

1. A method for writing file systems write data operations to a storage medium comprising the steps of:

storing a file systems write data operation in a first temporary data store;

mirroring the file systems write data operation in a second temporary data store; and

deleting the mirrored file systems write data operation from the second temporary data store in the case when the file systems write data operation is successfully written from the first temporary data store to the storage medium.

2. The method of claim 1, further comprising the step of:

writing the mirrored file systems write data operation from the second temporary data store to the storage medium in the case when the file systems write data operation is not successfully written from the first temporary data store.

3. The method of claim 1, further comprising the step of:

determining if the file systems write data operation is successfully written from the first temporary data store.

4. The method of claim 1, wherein said mirroring is performed one of concurrent with, during, or following said storing of the file systems write data operation in the first temporary data store.

5 The method of claim 1, further comprising the step of:  
sending a signal back to a source of the file system write data operation when it  
is determined that the file systems write data operation is successfully stored in the  
first and second temporary data stores.

6. The method of claim 1, wherein the storage medium is serviced by a  
plurality of servers, each server including the first and second temporary data stores  
and wherein:

    said storing includes storing the file systems write data operation in the first  
    temporary data store of the one of the plurality of servers;

    said mirroring includes mirroring the file systems write data operation being  
    stored in the first temporary data store of said one of the plurality of servers in the  
    second temporary data store of another of the plurality of servers; and

    said deleting includes deleting the mirrored file systems write data operation  
    from the second temporary data store of said another of the plurality of servers in the  
    case when the file systems write data operation is successfully written from the first  
    temporary data store of said one of the plurality of servers to the storage medium.

7. The method of claim 6, wherein said writing includes writing the  
mirrored file systems write data operation from the second temporary data store of  
said another of the plurality of the servers to the storage medium in the case when  
the file systems write data operation is not successfully written from the first  
temporary data store of said one of the plurality of servers.

TOP SECRET EEE EEE EEE EEE

8. The method of claim 6, wherein said determining includes determining if the file systems write data operation is successfully written from the first temporary data store of said one of the plurality of servers.

9. The method of claim 6, wherein said mirroring is performed one of concurrent with, during or following said storing of the file systems write data operation in the first temporary data store of said one of the plurality of servers.

10. The method of claim 6, wherein said writing the mirrored file systems write data operation from the second temporary data store of said another of the plurality of servers to the storage medium is done in the case when at least one of:

- (a) it is determined that said one of the plurality of servers is not operational;
- (b) it is determined that an operating system of said one of the plurality of servers is not operational; or
- (c) it is determined that an operating system I/O of said one of the plurality of servers is not operational.

11. A method for writing file systems write data operations to a storage medium comprising the steps of:

storing a file systems write data operation in a first temporary data store;  
mirroring the file systems write data operation stored in a second temporary data store;

determining if the file systems write data operation stored in the first temporary data store is successfully written to the storage medium;

deleting the file systems write data operation from the second temporary data store when it is determined that the file systems write data operation was successfully written from the first temporary data store to the storage medium; and

writing the mirrored file systems write data operation from the second temporary data store to the storage medium when it is determined that the file systems write data operation was not successfully written from the first temporary data store to the storage medium.

TOP SECRET//EGG//EYE//FO

12. The method of claim 11, wherein said mirroring is performed one of concurrent with, during, or following said storing of the file systems write data operation in the first temporary data store.

13. The method of claim 11, further comprising the step of: sending a signal back to a source of the file system write data operation when it is determined that the file systems write data operation is successfully stored in the first and second temporary data stores.

14. The method of claim 11, wherein the storage medium is serviced by a plurality of servers, each server including the first and second temporary data stores and wherein:

    said storing includes include storing the file systems write data operation in the first temporary data store of the one of the plurality of servers;

said mirroring includes mirroring the file systems write data operation being stored in the first temporary data store of said one of the plurality of servers in the second temporary data store of another of the plurality of servers;

said deleting includes deleting the mirrored file systems write data operation from the second temporary data store of said another of the plurality of servers in the case when the file systems write data operation is successfully written from the first temporary data store of said one of the plurality of servers to the storage medium; and

said writing includes writing the mirrored file systems write data operation from the second temporary data store of said another of the plurality of the servers to the storage medium in the case when the file systems write data operation is not successfully written from the first temporary data store of said one of the plurality of servers.

15. The method of claim 14, wherein said writing the mirrored file systems write data operation from the second temporary data store of said another of the plurality of servers to the storage medium is done in the case when at least one of:

- (a) it is determined that said one of the plurality of servers is not operational;
- (b) it is determined that an operating system of said one of the plurality of servers is not operational; or
- (c) it is determined that an operating system I/O of said one of the plurality of servers is not operational.

16. A method for writing file systems write data operations to a storage medium being serviced by a plurality of servers, each server including a first

TOP SECRET//EYES ONLY

temporary data store and a second temporary data store, said method comprising the steps of:

storing a file systems write data operation in the first temporary data store of one of the plurality of servers;

mirroring the file systems write data operation in the second temporary data store of another of the plurality of servers;

deleting the file systems write data operation from the second temporary data store of said another of the plurality of servers when it is determined that the file systems write data operation stored in the first temporary data store of said one of the plurality of servers was successfully written to the storage medium; and

writing the mirrored file systems write data operation in the second temporary data store of said another of the plurality of servers to the storage medium when it is determined that the file systems write data operation was not successfully written to the storage medium from the first temporary data store of said one of the plurality of servers.

17. The method of claim 5, further comprising the step of:

determining if the file systems write data operation is successfully written to the storage medium from the first temporary data store of said one of the plurality of servers.

18. The method of claim 17, wherein said mirroring is performed one of concurrent with, during, or following said storing of the file systems write data operation in the first temporary data store of said one of the plurality of servers.

19. The method of claim 16, further comprising the step of:  
sending a signal back to a source of the file system write data operation when it  
is determined that the file systems write data operation is successfully stored in the  
first temporary data store of said one of the plurality of servers and the mirrored file  
systems write data operation is stored in the second temporary data store of said  
another of said plurality of servers.

20. The method of claim 16, wherein said writing the mirrored file systems  
write data operation from the second temporary data store of said another of the  
plurality of servers to the storage medium is done in the case when at least one of:  
(a) it is determined that said one of the plurality of servers is not operational;  
(b) it is determined that an operating system of said one of the plurality of  
servers is not operational; or  
(c) it is determined that an operating system I/O of said one of the plurality of  
servers is not operational.

21. A system for writing file systems write data operations, comprising:  
a storage medium;  
a plurality of servers servicing the storage medium, each server including a first  
temporary data store and a second temporary data store; and  
a communications link, the communications link being configured and  
arranged so as to communicatively interconnect the first temporary data store of one  
of the plurality of servers to the second temporary data store of another of the

plurality of servers and to communicatively interconnect the first temporary data store of said another of the plurality of servers to the second temporary data store of said one of the plurality of servers.

22. The system for writing file systems write data operations of claim 21, wherein said communications link comprises a first and second communications interconnection, the first communications interconnection being configured and arranged so as to communicatively interconnect the first temporary data store of one of the plurality of servers to the second temporary data store of another of the plurality of servers and the second communications interconnection being configured and arranged to communicatively interconnect the first temporary data store of said another of the plurality of servers to the second temporary data store of said one of the plurality of servers.

23. The system for writing file systems write data operations of claim 22, wherein said first and second communications interconnects are each one of a fiber optical channel, a gigabit Ethernet and an infiniband.

24. The system for writing file systems write data operations of claim 21, wherein each server further includes a central processing unit and a program for execution on the central processing unit, said program including instructions and criteria for:

storing a file systems write data operation in the first temporary data store of said one of the plurality of servers; and

4325455262202020

communicating a copy of the file systems write data operation being stored in the first temporary data store of said one of the plurality of servers via the communications link to the second temporary data store of said another of the plurality of servers for storage of the copy therein;

25. The system for writing file systems write data operations of claim 24, wherein said program for execution on the central processing unit further includes instructions and criteria for:

deleting the copy of the file systems write data operation from the second temporary data store of said another of the plurality of servers in the case when the file systems write data operation is successfully written from the first temporary data store of said one of the plurality of servers to the storage medium; and

writing the copy of the file systems write data operation from the second temporary data store of said another of the plurality of the servers to the storage medium in the case when the file systems write data operation is not successfully written from the first temporary data store of said one of the plurality of servers.

26. The system for writing file systems write data operations of claim 25, wherein said program for execution on the central processing unit further includes instructions and criteria so that said communicating is performed one of concurrent with, during, or following said storing of the file systems write data operation to the first temporary data store of said one of the plurality of servers.

2025 RELEASE UNDER E.O. 14176

27. The system for writing file systems write data operations of claim 24, wherein said program for execution on the central processing unit further includes instructions and criteria for:

sending a signal back to a source of the file system write data operation when it is determined that the file systems write data operation is successfully stored in the first temporary data store of said one of the plurality of servers and the copy of the file systems write data operation is stored in the second temporary data store of said another of said plurality of servers.

28. The system for writing file systems write data operations of claim 24, wherein said program for execution on the central processing unit further includes instructions and criteria for performing said writing the copy of the file systems write data operation from the second temporary data store of said another of the plurality of servers to the storage medium in the case when at least one of:

- (a) it is determined that said one of the plurality of servers is not operational;
- (b) it is determined that an operating system of said one of the plurality of servers is not operational; or
- (c) it is determined that an operating system I/O of said one of the plurality of servers is not operational.

29. The system for writing file systems write data operations of claim 21, wherein:

file systems write data operations are sourced from one or more client computers of a computer network;

each of the plurality of servers includes a device that operably interconnects each of the plurality of servers to the computer network; and  
said program for execution on the central processing unit further includes instructions and criteria for receiving and processing each file systems write data operation being sourced from any one of the one or more client computers.

30. The system for writing file systems write data operations of claim 29, wherein said program for execution on the central processing unit further includes instructions and criteria for:

sending a signal back to said one of the one or more client computers sourcing a given file system write data operation when it is determined that the given file systems write data operation is successfully stored in the first temporary data store of said one of the plurality of servers and the copy of the given file systems write data operation is stored in the second temporary data store of said another of said plurality of servers.

31. A program for execution on a central processing unit of each server of a cluster of servers that service a storage medium, the cluster of servers for processing file systems write data operations to be written to the storage medium being and wherein each server includes a first and a second temporary data store, said program comprising instructions and criteria for:

storing the file systems write data operation in the first temporary data store of one of the servers of the cluster;

mirroring the file systems write data operation being stored in the first temporary data store of said one of the servers of the cluster in the second temporary data store of another of the servers of the cluster;

deleting the mirrored file systems write data operation from the second temporary data store of said another of the servers of the cluster in the case when the file systems write data operation is successfully written from the first temporary data store of said one of servers of the cluster to the storage medium; and

writing the mirrored file systems write data operation from the second temporary data store of said another of the servers of the cluster to the storage medium in the case when the file systems write data operation is not successfully written from the first temporary data store of said one of the servers of the cluster.

32. The program for execution on a central processing unit of each server of a cluster of servers of claim 31, wherein the program further includes instruction and criteria for:

monitoring an operational status of each of the servers of the cluster; and causing the writing of mirrored file systems write data operation from the second temporary data store of said another of the servers of the cluster to the storage medium, when it is determined from said monitoring that said one of the clusters of the server is not operational.

33. The program for execution on a central processing unit of each server of a cluster of servers of claim 31, wherein a plurality of mirrored file systems write data operations are stored in the second temporary data store of said another of the

TOP SECRET//COMINT

servers of the cluster and wherein said writing includes writing all of the plurality of mirrored file systems write data operations.

34. A method for writing file systems write data operations to a storage medium being service by a plurality of servers, comprising the steps of:

storing the file systems write data operation in the first temporary data store of one server;

mirroring the file systems write data operation being stored in the first temporary data store of said one server in the second temporary data store of another server;

monitoring the operational status of the each server; and

writing the mirrored file systems write data operation from the second temporary data store of said another server to the storage medium in the case when said monitoring determines that said one server is not operational.

35. The method for writing file systems write data operations to a storage medium of claim 34, wherein a plurality of mirrored file systems write data operations are stored in the second temporary data store and wherein said writing includes writing all of the plurality of mirrored file systems write data operations from the second temporary data store.

36. The method for writing file systems write data operations to a storage medium of claim 35, wherein said writing of all of the plurality of mirrored file system

write data operations are completed before accepting any new file system write data operations for writing to the storage medium.

37. The method for writing file systems write data operations to a storage medium of claim 35, further comprising:

stopping said mirroring of file systems write data operations in the case when said one server is not operational; and

restarting said mirroring of file systems write data operations in the case when said one server is returned to operation.

38. A method for writing file systems write data operations to a storage medium being serviced by a plurality of servers, each server including a first temporary data store and a second temporary data store, and wherein each server is configured and arranged so that one server has ownership over a first portion of the storage medium and another server has ownership of a second portion of the storage medium, said method comprising the steps of:

storing a given file systems write data operation in the first temporary data store of one of the plurality of servers, the server having ownership over that portion of the storage medium the given file systems write data operation is to be written to;

copying the given file systems write data operation being stored in the first temporary data store in the second temporary data store of the other of the plurality of servers, the server that does not have ownership over that portion of the storage medium the given file systems write data operation is to be written to;

deleting the copy of the given file systems write data operation from the second temporary data store when it is determined that the given file systems write data operation stored in the first temporary data store is successfully written to the storage medium; and

writing the copy of the given file systems write data operation in the second temporary data store to the storage medium when it is determined that the given file systems write data operation was not successfully written to the storage medium from the first temporary data store.

40. The method for writing file systems write data operations to a storage medium of claim 38, further comprising the steps of:

monitoring the operational status of each of the plurality of servers;

determining which server is inoperable and a particular portion of the storage medium the inoperable server has ownership over;

assigning ownership of the particular portion of the storage medium to the operational server; and

causing the writing of the copy of the given file systems write data operation from the second temporary data store of the operational server to the particular portion of the storage medium.

40. The method for writing file systems write data operations to a storage medium of claim 39, wherein a plurality of given file systems write data operations are stored in the second temporary data store and wherein said writing includes writing

all of the plurality of given file systems write data operations from the second temporary data store.

41. The method for writing file systems write data operations to a storage medium of claim 40, wherein said writing of all of the plurality of mirrored file system write data operations is accomplished before accepting any new file system write data operations for writing to the storage medium.

42. The method for writing file systems write data operations to a storage medium of claim 39, further comprising the steps of:

stopping said mirroring of file systems write data operations when an there is an inoperable server; and

restarting said mirroring of file systems write data operations when the inoperable server is returned to operation.

TOP SECRET//  
REF ID: A6560